

HPC at ESRL for 2011

An Overview

February 3, 2011



Overview

- New ESRL HPC Site Lead
- Recovery Act Funding and NOAA HPC
- Upcoming Network Changes/NWAVE
- HFIP Upgrade



GPU Technical Seminar



New ESRL HPC Site Lead

- We now have a permanent replacement for the ESRL HPC Site Lead: Forrest Hobbs
- Forrest has been site lead for the integrator for over 9 years



American Reinvestment and Recovery Act

April 9th, 2009 it was announced NOAA had received funds to expand HPC

century. The Act provided a total of \$170 million in funds for NOAA programs for climate modeling activities, including procurement of the supercomputer and research into climate change.

Centralization of HPC resources is now upon us!



Where did the money go?

- Through a partnership with Oak Ridge National Laboratory, \$70m went to build a machine for Climate Computing
 - First phase completed last year – Gaea
 - System will eventually be over 1 Pflops (late 2011)
 - 120,000 cores, Cray XE6
 - Majority of cycles for GFDL/Climate users, but there are cycles available for other NOAA users



Where did the money go? (Cont)

- \$70m dedicated to support research in weather modeling
 - New NOAA-wide data center
 - NWAWE network
 - New HPC system and HSMS
 - Procurement ongoing now

**New Site: National Environmental Security
Computing Center (NESCC)**



Timeline

- April 2009 – ARRA funds announced
- Summer 2009 – Winter 2010
 - Define requirements for systems
- March 2010 – Award new integration contract to Computer Sciences Corporation
- August 2010 – Release RFI for new systems
- October 2010 – Ribbon-Cutting Ceremony on Vertex Center, Home of the NESCC
- February – CSC to award contract to HSMS vendor
- March 2011– CSC to award contract to HPC system vendor
- June 2011 – New HSMS to go online
- October 2011 – Compute system to go online



Vertex Center - Home of NESCC



Vertex Center Details

- Located in Fairmont, WV
- 20 year lease
- 19,000 sq. ft.
 - Large enough to house several Pflops systems
- Half of room saved future NOAA HPC needs
- Other NOAA organizations will eventually relocate here as well



NESCC Details

- Procurement for NESCC on going now
- HPC system should be at least 325 Tflops
- HSMS should hold at least 50 Pbytes
 - 80-90% for NCEP R&D/Operations needs
- Majority of resources will be for existing work coming from ESRL and NCEP

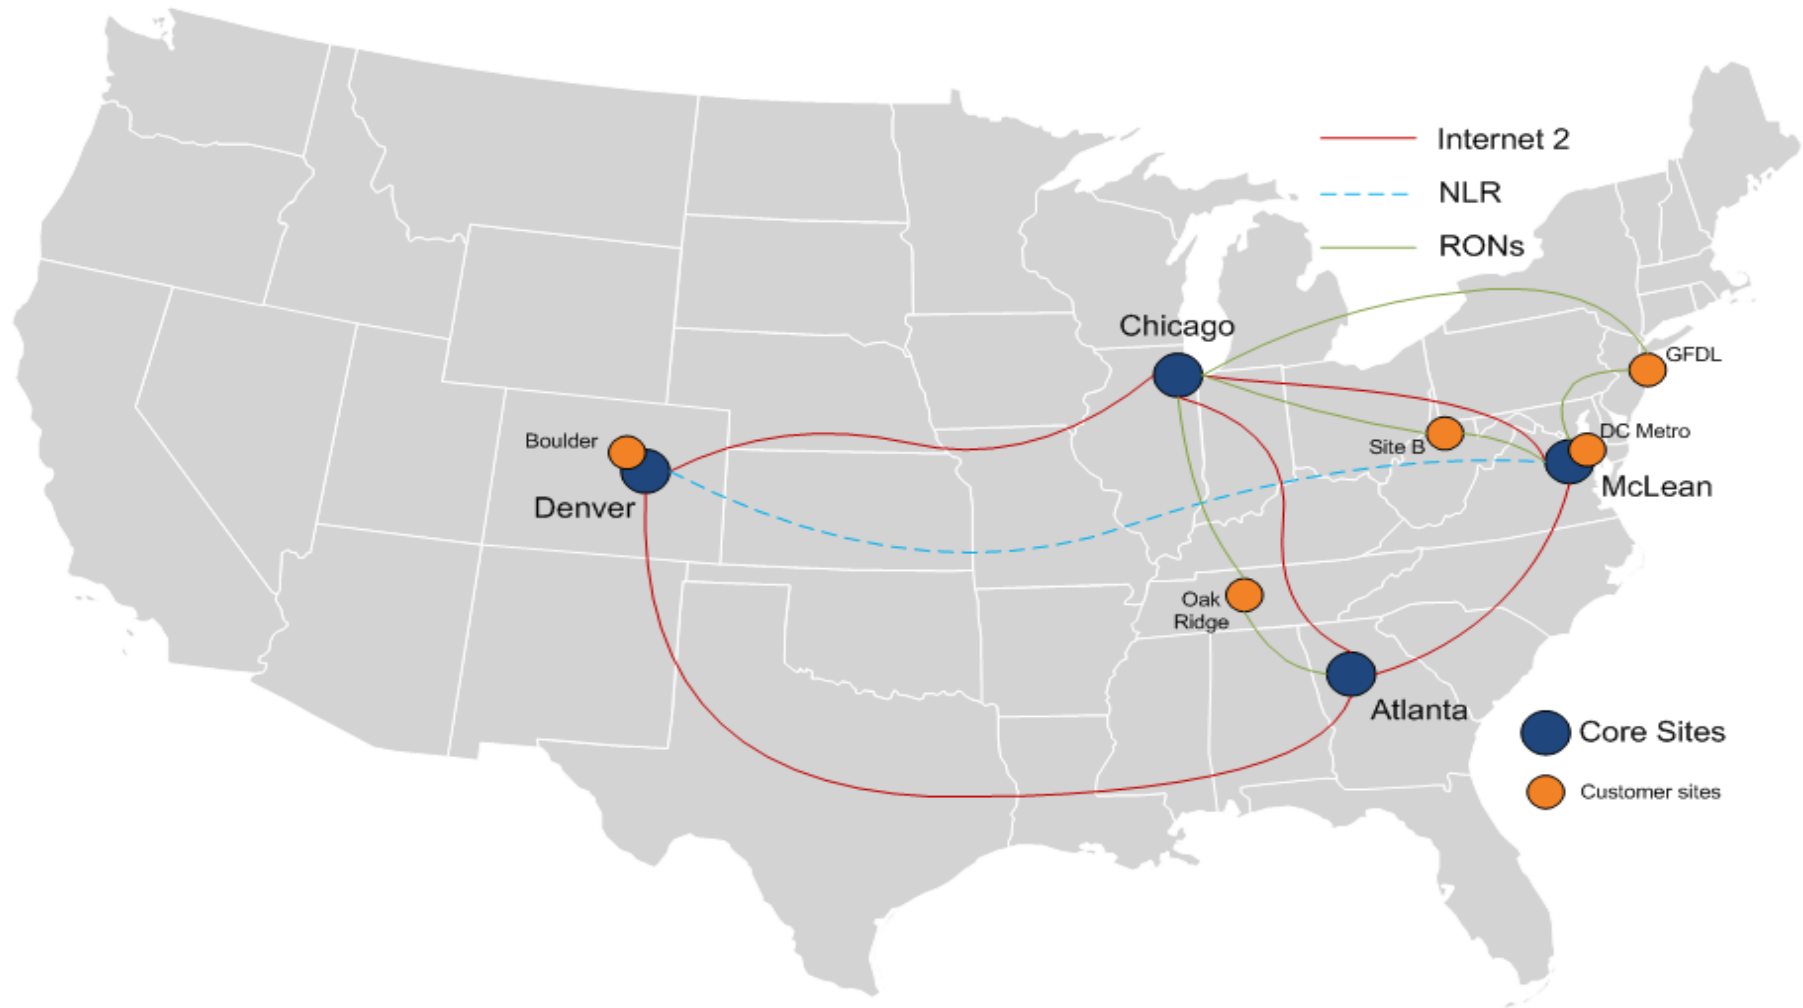


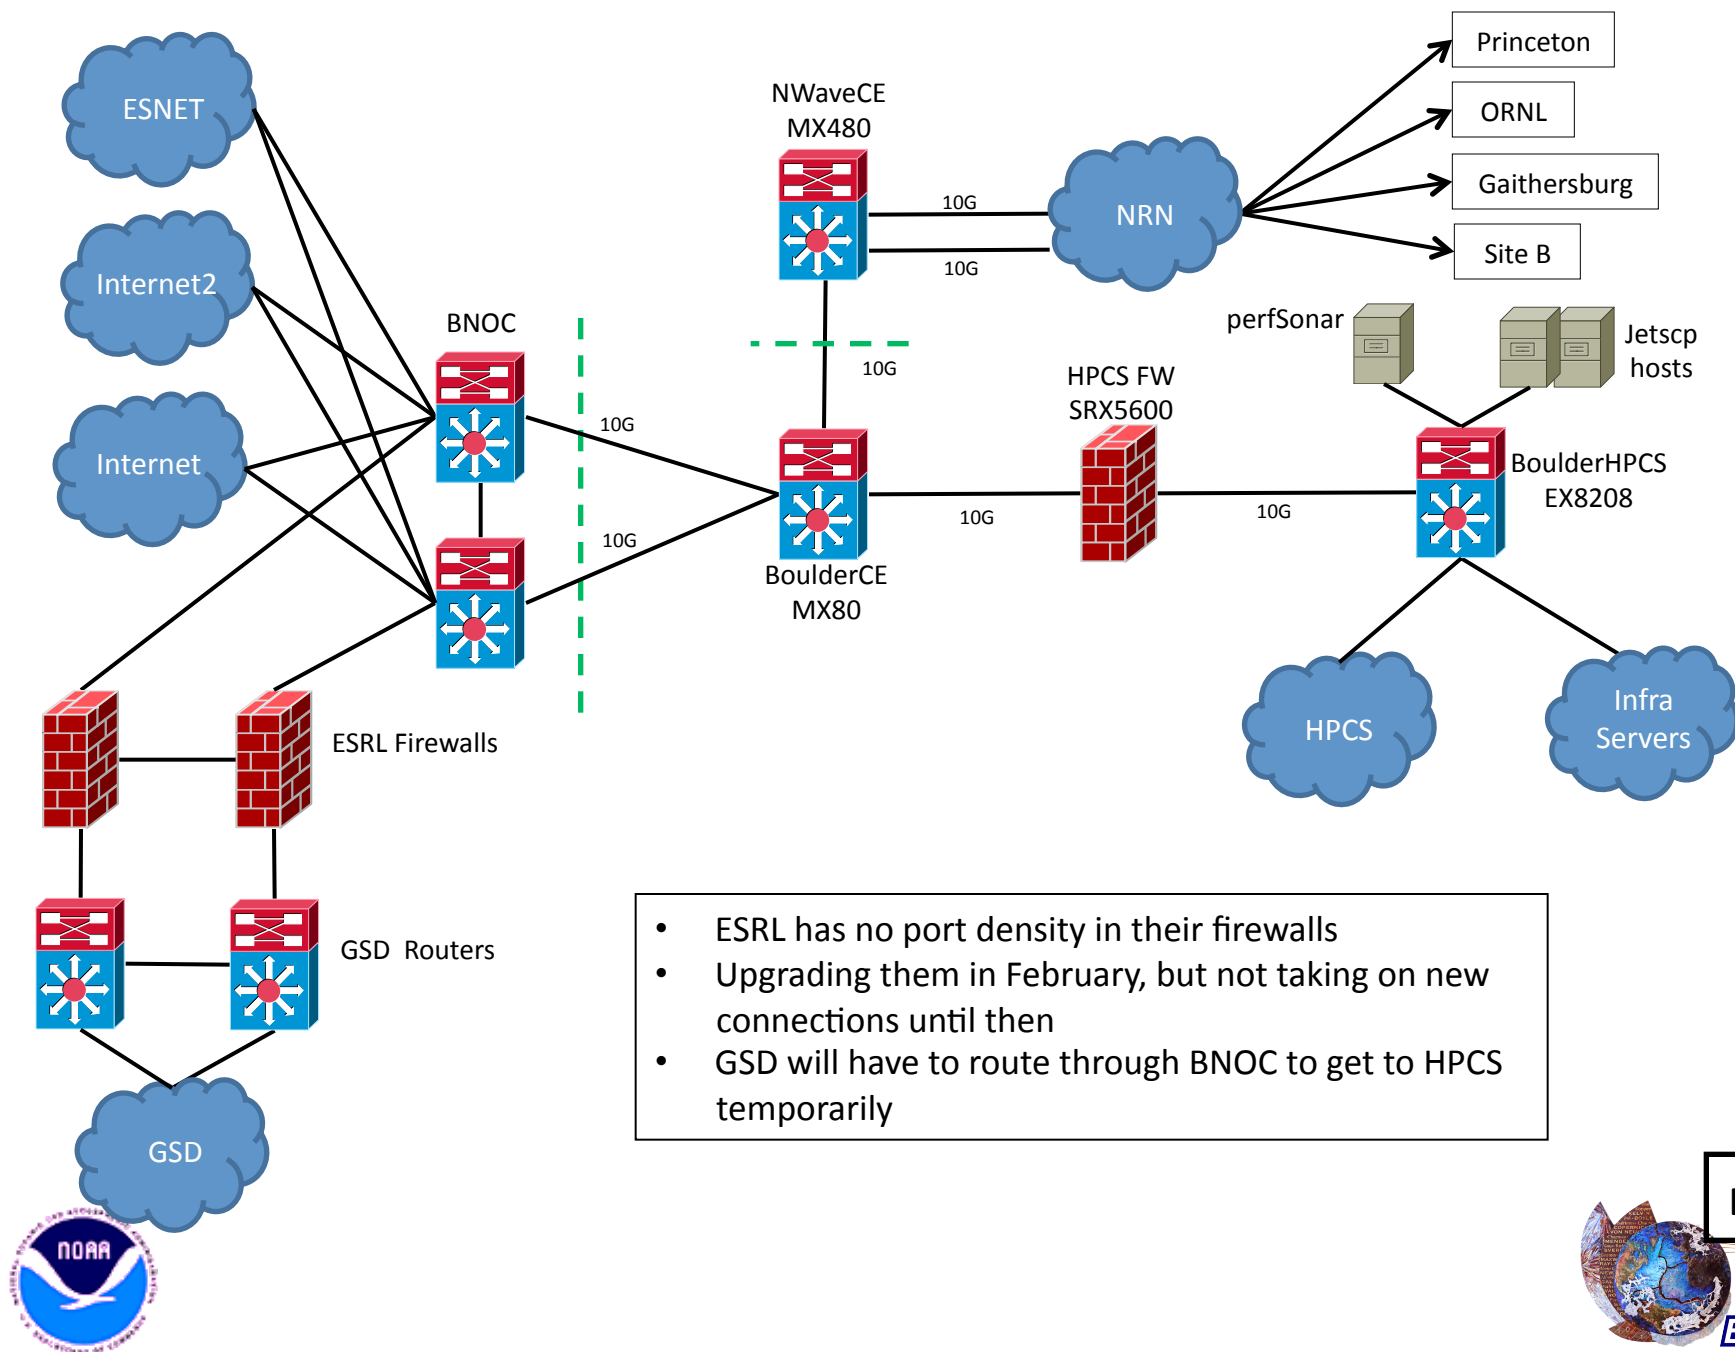
NESCC

- Leslie Hart is the Site Lead for NESCC
 - No, he isn't moving
- Forrest Hobbs
 - Site Lead for ESRL
 - Deputy Site Lead for NESCC
- Support staff will be distributed at the existing major HPC sites (ESRL and NCEP)



NWAVE





- ESRL has no port density in their firewalls
- Upgrading them in February, but not taking on new connections until then
- GSD will have to route through BNOc to get to HPCS temporarily

NWAVE in Boulder

- New equipment will provide 10 Gbit/s networking NWAVE and the Internet
 - Performance can be easily upgraded with small upgrades to new hardware
- Improvements will be for external users, not just NWAVE sites
 - Early testing of new Internet connection demonstrates over 100 MB/s on a single Gbit/s link (over 80% peak)
 - We are working to tune the 10 Gbit/s link



Future of RDHPCS Resources

- What does this mean for the RDHPCS users in Boulder?
 - HSMS system will be located in WV, data will be written there from Boulder
 - Existing RDHPCS workloads will be migrated there
 - Technical requirements stated existing Boulder resources must be operational 6 months after acceptance of NESCC
 - Around March, 2012



HSMS Usage in Boulder

- All RDHPCS data will be written to the new HSMS
 - Should start around June 2nd
- Good news
 - Much more capacity than before
 - Will be near the new HPC resources which will be a benefit when the new system comes on line
- Not so good news
 - We cannot direct mount HSMS on our front ends
 - Will need to use a mssGet/mssPut like script for moving data
 - State of current HSMS for GSD/ESRL unknown



How will I use the remote system?

- Logins will start like they do now, authenticate into Jet
- Additional hop to NESCC
 - This may become transparent
- Data
 - There will be a way to move data from external sites directly to NESCC filesystems as done with jetscp – Mechanism may change
 - There will be optimized ways to move data between Jet and NESCC, scripts will be provided
 - Data from /public and other real-time data feeds will be available at NESCC
- Batch system
 - SGE will not be the batch system at NESCC
 - Changes to scripts will be required
 - Changes to scripts can be made portable back to any Jet resources
- Compilers, batch system, and architecture not yet specified

We do not have all the answers yet. We intend to hold more frequent user meetings this spring and summer as details are worked out.



HFIP for 2011

- HFIP should be getting an upgrade before hurricane season in 2011
 - Same hardware as tJet (4000-5000 cores)
 - Additional disk
 - We are working the details on this now
 - Current holdup – No federal budget approved
- As to which HSMS HFIP will be storing their long term data is to be determined



Upcoming Changes

- New Infrastructure
 - Infrastructure connects Jet to the outside world
 - Provides standard services of any Internet connected system
 - Necessary for NWAWE connectivity
- Individual service outages between Feb 3rd and 17th
- Downtime 3rd week of Feb, 8am-6pm to complete move
 - Exact date to be determined
 - Users should not see a difference except network performance
 - SSH connectivity from Jet to GSD/ESRL will be broken
 - Wide open access is not secure
 - Exceptions can be requested for unattended data transfers back into the GSD/ESRL environment



Questions???

